

APPLICATION OF

SYED NOMAN KAZMI,
VINAY POLAVARAPU
and
WILLIAM TURNER

FOR LETTERS PATENT OF THE UNITED STATES

METHOD AND SYSTEM FOR MANAGING DIGITAL
CONTENT INCLUDING STREAMING MEDIA

James J. DeCarlo
Registration No. 36,120
Attorney for Applicants
STROOCK & STROOCK & LAVAN LLP
180 Maiden Lane
New York, New York 10038
(212) 806-5400

Atty. Docket No.: 694231/0086

METHOD AND SYSTEM FOR MANAGING DIGITAL CONTENT, INCLUDING STREAMING MEDIA

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation-in-part of, United States Application No. 10/466,132 filed on July 9, 2003 which is based on International Application No. PCT/US02/01840 filed on January 18, 2002 for METHOD AND SYSTEM FOR MANAGING DIGITAL CONTENT, INCLUDING STREAMING MEDIA which claims the benefit of, United States Provisional Application No. 60/263,058, filed on January 18, 2001, for METHOD AND SYSTEM FOR MANAGING STREAMING MEDIA. This application also claims the benefit of United States Provisional Application No. 60/491,802, filed July 31, 2003 for METHOD AND SYSTEM FOR MANAGING DIGITAL CONTENT, INCLUDING STREAMING MEDIA, all of which are hereby incorporated by reference.

BACKGROUND OF INVENTION

1. Field of the Invention

[0002] The present invention relates generally to a method and system for providing digital content via the Internet and, more particularly, to a method and system for allowing clients to up-load, manage, and deliver streaming media content via the Internet.

2. Description of Related Art

[0003] With the advent of the Internet and the World Wide Web, an industry has developed around the delivery of digital content, such as streaming media content. By way of example, streaming media may be used for any of a number of purposes, including entertainment, distance learning and corporate purposes. Entertainment companies stream movies and sporting events, distance learning companies stream educational content, and corporations stream training materials.

[0004] Although some streaming media content providers may have relatively few items of content to provide, some content providers have hundreds, even thousands of content files. Storing and streaming this number of content files can be costly. Furthermore, streaming content requires a level of technical expertise often not found in companies focusing on creating content. Consequently, content providers have turned to content management service providers to store and stream content on behalf of the content providers.

[0005] As more content providers turn to service providers for their streaming technology needs, service providers must manage more client accounts and greater numbers of content files. Furthermore, to maintain existing content provider clients and attract new clients, service providers must provide a content management system that not only is capable of organizing large numbers of content files, but also is easy for the content providers to use. Accordingly, there exists a need for a content management system that allows clients to easily up-load, manage, and deliver streaming media content via the Internet.

3. Summary of the Invention

[0006] The present invention solves the foregoing and other needs. In certain embodiments, a method and system is provided for allowing client content providers to upload, manage and deliver streaming media and other digital content. The system according to one embodiment receives digital content from the client, assigns a stream identifier (ID) to the content and stores the content. The client is given a playlist uniform resource locator (URL) for publishing on its web site, the URL including the stream ID. Activation of the URL by an end user causes the stream to be served to the end user, without the client receiving or providing an indication of the specifics of where the content was stored.

[0007] Another embodiment of the present invention permits clients to actively manage their content, including defining logical folders and subfolders containing item(s) of content; defining logical stream groups, containing items of content.

[0008] A system according to one embodiment of the present invention couples one or more media servers to a storage server. While the storage server has stored a copy of the content, the media servers coupled thereto do not. In operation, the media servers perform a read of the contents stored at the storage server when requested by an end user.

[0009] Another embodiment includes a system and method for facilitating access by users to digital content of one or more clients, the clients having client identifying information. Such a system includes one or more remote sites, each comprising one or more first servers configured to receive digital content from clients, the digital content to be available to users via the remote sites. The system also includes a central site geographically remote from the remote sites, the central site comprising one or more ingest queue servers to direct transfer of digital content to one or more of the first servers based on the client identifying information.

BRIEF DESCRIPTION OF THE DRAWINGS

[0010] The following drawing figures, which form a part of this application, are illustrative of embodiments of the present invention and are not meant to limit the scope of the invention in any manner, which scope shall be based on the claims appended hereto.

[0011] Figure 1 is a schematic illustrating the system architecture of one embodiment of the present invention;

[0012] Figure 2 is a schematic illustrating the database of one embodiment of the present invention;

[0013] Figures 3a-j are web pages of the Content Management client-side web site according to one embodiment of the present invention;

[0014] Figures 4a - b are flowcharts illustrating processes of up-loading content files into the content management system according to one embodiment of the present invention;

[0015] Figure 5 is a flowchart illustrating the process of creating a playlist according to one embodiment of the present invention;

[0016] Figure 6 is a schematic illustrating the system architecture of an alternate embodiment of the present invention;

[0017] Figure 7 is a schematic illustrating a workflow of a content management system according to the alternate embodiment of Figure 6;

[0018] Figures 8a and 8b represent a schematic illustrating the database of the embodiment of Figure 6;

[0019] Figure 9 is a schematic illustrating the system architecture of an alternate embodiment of the present invention; and

[0020] Figure 10 is a schematic illustrating a workflow of the system according to the alternate embodiment of Figure 9.

DETAILED DESCRIPTION OF CERTAIN EMBODIMENTS

[0021] Certain embodiments of the present invention will now be discussed with reference to the aforementioned figures, wherein like reference numerals refer to like components. Turning first to the schematic of Figure 1, a system 100 according to one embodiment of the present invention is shown. In general, the system 100 allow clients 102 to up-load streaming media content to the system 100, manage such content, and make the content available to end-users 104 via a web site on the Internet.

[0022] As will become apparent from the following discussion, each of the clients 102 and end-users 104 has a processor, such as a personal computer (PC), web enabled cellular telephone or personal digital assistant (PDA) and the like, coupled to the Internet by any one of a number of known manners. Furthermore, each client 102 preferably includes an Internet browser, such as that offered by Microsoft Corporation under the tradename INTERNET EXPLORER or that offered by Netscape Corp. under the tradename

NETSCAPE NAVIGATOR, as well as file transfer protocol (FTP) client software for interacting with various components of the system 100. Additionally, each end-user processor preferably includes an Internet browser, such as INTERNET EXPLORER or NETSCAPE NAVIGATOR, and a media player, such as that offered by Microsoft Corporation under the tradename WINDOWS MEDIA PLAYER or that offered by Real Networks Corp. under the tradename REALPLAYER. It is to be understood that although the present embodiment is described in terms of Windows Media content and Real Media content, it is within the scope of the present invention to utilize any media format heretofore or hereafter known. Furthermore, it is to be understood that although the present embodiment is described in the context of streaming media, the present invention is applicable to digital content other than streaming media.

[0023] As discussed in greater detail below, the client 102 up-loads streaming media content to the repository or storage server (and associated storage) 106 either directly, via an HTTP up-load, or via an FTP ingest server (and associated storage) 108. Once the client 102 has up-loaded its streaming media content, the client may manage its content via web pages on a content management web site provided by a client-side web server 110. As discussed in greater detail with reference to Figure 3a-j, the web site serves as an interface through which the client 102 may log into its account and select any of the following exemplary functions: up-load new content; search the client's existing up-loaded content; create and edit playlists; browse the client's contents stored on the FTP server 108; and perform a batch up-load of content from the FTP server 108 to the repository server 106. Additionally, the web site serves as an interface for the client 102 to perform various administrative functions, such as setting and changing any of the client-defined account and file information described below. In alternate embodiments, clients 102 may manage their accounts programmatically, via scripts running on the web server 110 or other system service described below.

[0024] In various different embodiments, client 102 may provide various types of content through an application, such as a web application that creates and uploads content into the system, which content eventually gets provided to individual users of a web site, an application for creating a webcast and the like. This content can be any type, including, for example, music videos, television programs, distance learning, video conferences and other web broadcasts, which can be uploaded via FTP, HTTP or other ingest mechanisms as described herein. Alternatively, client 102 can refer to certain other web properties or pages, such as those directed to certain vertical markets like News, Sports, Financial, Personals, Weather and other information. In yet another embodiment, the content can be uploaded by internal users of the system who upload files on behalf of clients.

[0025] The system further includes a file management server 112, administrative staff terminals 114, and a script processor 116. In general, the file management server 112 controls the activities of the various other components to which it is coupled and the movement of content within the system. The administrative terminals 114, in turn, are coupled to the file management server 112 and database server 118 to allow the administrative staff of the content management service provider maintaining the content management system 100 to control and monitor the system 100. The script processor 116 works in conjunction with the file management server 112 to control operation of the system 100 by running various software scripts and/or components performing most of the functionality described herein. As described in greater detail below, the script processor 116 includes various software modules, including a task scheduler and program scripts and objects for batch up-loads of content files and for recovery of deleted content files. It is to be understood that such task scheduling software may be obtained by any of a number of third-party vendors. Similarly, based upon the description of the program scripts and objects

herein, the program scripts and objects may be written in any programming language heretofore or hereafter known, such as PERL, Visual Basic, JavaScript, C++, and the like.

[0026] The content management system 100 also includes a Content Management (CM) Database 118. As described in greater detail below with reference to Figure 2, the CM Database 118 includes numerous relational databases or tables in a database. In general, the CM Database 118 preferably includes account information, which identifies each client's account, stream information, which identifies and describes each item of streaming content within a client's account, playlist information, which identifies and describes each client's playlist, batch information, which defines the status of each client's batch up-load requests, and storage location information, which tracks the storage of content on the repository 106 and streaming media 120 servers.

[0027] Under direction of the file management server 112, the up-loaded content is eventually transferred from the repository server 106 to one or more streaming media server (and associate storage) 120, where it is available for streaming via the Internet. As described in greater detail below, the current embodiment utilizes a playlist server 122 for dynamically generating a playlist metafile (such as ASX, RAM, SMIL, and RPM files, to name a few) to be provided to the end-user's windows media player. In addition to the archived content located on the streaming media server storage 120, the present environment provides for streaming of live content acquired via encoders 124 coupled to the streaming media servers 120.

[0028] As illustrated in Figure 1, while the clients 102 and end-users 104 are coupled to each other and to the content management system 100 via the Internet, the service provider's components are coupled to each other via a local area network (LAN). More specifically, the repository 106, FTP server 108, client web server 110, file management 112, administrative terminals 114, script processor 116, CM Database 118, streaming media server

120, and playlist server 122 are all in communication via a secure LAN. In one embodiment, the system provides additional security by using Secure Sockets Layer (SSL) connections, which is well known in the art. In alternate embodiments, components of the system are coupled differently, for example, each coupled directly to the Internet, coupled via a wide area network (WAN) and the like.

[0029] In general, the division of functionality among the web server 110, file management server 112, playlist server 122 and script processor 116 described herein is exemplary, as the functions may be segregated and grouped differently. For example, it is to be understood that although the present embodiment utilizes client web servers 110 that are separate from the playlist server 122, in alternate embodiments the functionality of the playlist server 122 described herein may be implemented in software residing on the web server 110, thereby obviating the need for a separate playlist server.

Content Management Database 118

[0030] The Content Management (CM) Database 118 will now be described in greater detail with reference to Figure 2, and continuing reference to Figure 1. The CM Database 118 includes several logically discrete tables of information, each of which is described below. As will be appreciated by one skilled in the art, the following logical arrangement of information in tables is exemplary, and other arrangements are within the scope of the present invention; for example; those with fewer or more tables and/or fewer or more fields. It is also to be understood that although the CM Database 118 is described as a central database, it is within the scope of the present invention to distribute the contents of the CM Database 118 throughout the system.

[0031] The CM Database 118 includes a CM Accounts Table 210. In general, the CM Accounts Table 210 includes records for each client account, as identified by a CM account ID. Each such record includes account identifying information, such as account

name, account directory, username, password, contact information (such as contact name, e-mail address, and telephone number), default title and author (which are preferences used to identify items of content in the event no other title or author information is provided by the client 102), parent account ID (in the event the current account is related as a child to another account), an indication of whether or not the client has an FTP account, and maximum allowable amount size of content (e.g., in Kilo-bytes). It is within the scope of the present invention to implement security features to maintain the integrity of the content, such as those features and methods described in applicant's co-pending International Application No. PCT/US01/46726, filed November 5, 2001, for SYSTEM AND METHOD FOR CONTROLLING ACCESS TO DIGITAL CONTENT, INCLUDING STREAMING MEDIA, and International Serial No. PCT/US01/18324, filed June 6, 2001, for SYSTEM AND METHOD FOR CONTROLLING ACCESS TO DIGITAL CONTENT, INCLUDING STREAMING MEDIA, which are hereby incorporated by reference. As such, the table includes security related information, including default preferences indicating whether or not the content is secure (allows security), and if so, a cryptographic key used for accessing such content (default sec key) and a time period for which the key will be valid (default sec. interval).

[0032] The CM Accounts Table 210 also includes a field for a "site id," which uniquely identifies a site consisting of a group of one or more FTP servers 108 and those repository servers 106 and streaming media servers 120 to which the content on the FTP servers 108 can be up-loaded. Two representative sites are illustrated in Figure 1, although the system 100 could have any number of sites. Although not required in the present invention, each client account is associated with a site located in the client's geography or the geography where the endorsers 104 likely to access the client's content are located so to speed delivery of the streaming media content. Furthermore, it is to be understood that logically

grouping the FTP server 108, repository servers 106 and streaming media servers 120 into sites is not required. As noted below, much of the account identifying information of the present embodiment is provided by the client 102 during the registration process.

[0033] The CM Database 118 further includes a series of tables containing content – identifying information. More specifically, the CM Assets Table 212 includes records, each of which identifies an item of streaming content by a stream ID. Furthermore, each record includes the CM account ID, creation date of the item of content, description of the content, an identification of whether the content is audio and or video, the platform or format to which the content relates, the date on which the content was last modified, any code necessary for viewing of the content, the length and size of the content, the expiration date (if any) of the content, the individual that up-loaded the content and the date on which the content was processed into the system 100. It is to be understood that each record in the CM assets Table 212 correlates an item of content, as identified by the stream ID, to a particular client account, by reference to the CM account ID.

[0034] It is to be understood that the present invention is applicable to any platform and format now known or hereinafter developed including for example, MP4, Flash COM, Real Media, MP3, MP3PRO, QuickTime, Windows Media, a GIF Image, a JPEG Image, or XML.

[0035] The Streams2 Table 214 also includes records identifying each item of content, as identified by the stream ID. Fields contained in the Streams2 Table 214 include the stream type, such as windows media player or real media player, title of the content, author of the content, status of the content, copyright notice for the content, URL prefix, bite rate of the content, file name of the content (as provided by the client 102), stream format, such as windows media player or real media player, ad tag (which specified where an

advertisement is to be inserted), start time (if the Stream is a portion of the file), duration, and expiration date.

[0036] As illustrated in Figure 2, the CM Assets Table 212 has three additional tables related thereto: the CM Stream Keywords Table 216, CM Assets Locations Table 218 and CM Archive Volume Table 220. The CM Stream Keywords Table 216 includes keywords, as provided by the client 102, for each item of content as identified by the stream ID.

[0037] The CM Assets Locations Table 218, which is also related to the CM Assets Table 212, includes records identified by location ID. The location ID is a unique identifier for each portion of storage pertaining an item of content. As such, location IDs may refer to either a repository server 106 or a streaming media server 120. A given content file may be stored at multiple location ids, for example one on a media server 120 and one on a repository server 106. More specifically, each record in the CM Asset Locations Table 218 includes a CM volume ID, for identifying the portion of the server on which the associated content resides; stream ID, for identifying the item of content stored at the location ID; and status of the item of content. The status of the file corresponding to each location is determined based on the following table.

<u>Status</u>	<u>Meaning</u>	<u>Applicable to file located on:</u>
Available	File is present on a media server and is available for streaming	Streaming server
Back up	File has been copied to a streaming media server	Repository server
Copying	File is being copied to a repository server or media server	Repository server; media server
Deleting	File is being deleted	Repository server; media server
Failed	File up-load or copy has failed	Repository server; media server
Moving	File is being moved from the current server to another server of the same type	Repository server; media server
New	File has recently been up-loaded and not copied	Repository server

[0038] The CM Archive Volume Table 220 is also related to the CM Asset Table 212 and, more particularly, to the CM Asset Locations Table 218. CM Archive Volume Table 220 includes records for each CM volume ID. Each record includes: volume type, which indicates whether the volume ID pertains to a repository server 106, streaming media server 120, FTP server 108, script processing server 116 or web server 110; archive server name, indicating the LAN name of the particular repository server 106 or streaming media server 120, as the case may be, containing the volume ID; the server's DNS address; the platform for the content, such as Window's Media or Real Media; the status of the content associated with the volume id; the maximum storage capacity for the server containing the volume id; and the site id of the server containing the volume id.

[0039] In one embodiment, the following values represent the following volume types:

Value	Volume Type
F	FTP Server
P	Processing Server
S	Storage Server
W	Web Server

[0040] Also associated with each client account identified in the CM Accounts Table 210 is client playlist information. Such playlist information is contained within the CM Playlists Table 222, Content Group (CG) Table 224 and CG Streams Table 226. In general, the CM Playlist Table 222 includes records identifying each client playlist, as identified by a content group ID. Each record further includes the CM account ID and the sort order, which indicates the ordering of the items of streaming content associated with the particular playlist.

[0041] Content Group Table 224 also includes records for each playlist, and identified by the playlist or content group (CG) ID. In general, each record contains the playlist details, including the CG ad tag, the type of playlist, the CG description, the CG format, the CG label, the meta-url for the playlist, CG notes, and an indication as to whether or not the content group is shared.

[0042] The playlist entries are identified in a CG Stream Table 226. More specifically, the CG Stream Table 226 includes records identifying each item of content as part of a content group and further includes an indication of the order of the particular item of content within the content group. As such, each record in the table includes the content group ID, stream ID, and sort order.

[0043] The CM Database 118 also includes a table that indicates its status of client-requested batch up-loads. More specifically, the CM Batch Jobs Table 228 includes records identifying the status of each requested batch job, as identified by a CM batch job ID. Each record in the CM Batch Job Table 228 includes the CM account ID, identifying the account to which the batch job relates, the file name of the batch file (filename), the time which the

client 102 submitted the batch request (submit time), the time at which the process was started (process start time), the time at which the batch process ended (process end time), and an indication as to whether or not the batch file was removed from the client's account (file removed). Because each record in the CM Batch Jobs Table 228 includes the CM account ID, each such record is related to a record in the CM Accounts Table 210.

[0044] The CM Database also includes a Stream-Servers Table 230. The records in the Stream-Servers Table 230 correlate each file, as identified by its stream ID, with the hostname (or domain name server (DNS) address) of the streaming server 120 on which it resides. While each stream ID will only have one record, a particular file will have as many records as it has copies residing on different streaming servers 120.

[0045] The Servers2 Table 232 includes a record for each streaming server 120, as identified by its server address, setting forth various metrics for the server 120, such as the weight or relative number of streams being served to end users 104.

[0046] The CM Sites Table 234 contains a record for each site, as identified by its site id. More specifically, each record identifies the name of the site, the FTP server uniform resource locator (URL) associated with the site and the processor URL associated with the site. In the present embodiment, each site only includes one FTP server 108.

[0047] As illustrated in Figure 2, the CM Database also includes tables relating to the deletion and recovery of files, identifying each client's service level and categorizing the clients' accounts.

[0048] The Account Servicelevels Table 236 stores a description of each client's service level. As illustrated, this table 236 is keyed to the Accounts table 210 by virtue of the account service level ID (acct srvlevel id) field. Each service level generally corresponds to a different set of options, functions and/or capabilities offered to the client, for example, the amount of storage allowed, number of uploads, the number of sites provided, and the like.

[0049] The Categories Table 238 stores records identifying the category to which a particular account relates. More specifically, each record corresponds to a particular category, as identified by category ID. In general, a category is a grouping of accounts, for example, of all accounts related to a particular parent account. Thus, each record identifies the parent ID, name, and description. As illustrated, records in the Categories Table 238 are keyed to accounts in the Accounts table 210 by the category ID field.

[0050] The CM database 118 also tracks deleted accounts and deleted content. To this end, the Deleted Accounts Table 240 keeps a record for each account, identified by account ID, that is removed from the system 100. By keeping such records, the system administrator is able to provide accurate record keeping and report generation.

[0051] The Deferred Delete Table 242, Deleted Assets Table 244 and the Recovery Jobs Table 246 are used in the deletion of content. More specifically, a Stream may be deleted either upon request of the client 102 or based on the content expiring. As noted above, the Streams2 Table 214 includes an expiration date field for each Stream ID. A script running on the system periodically scans the Streams2 Table 214 for records having an expiration date equal to or greater than (i.e., after) the then current time. For each record containing such an expiration date, the system proceeds to delete the content. The process of deletion of the present embodiment includes changing the status field in the Asset Locations Table 218 to expired and, in the Streams-Servers Table 230 changing the hostname for the stream being deleted to another system server that hosts a default message, indicating to the end user 104 that the requested content is not available. The system also replaces the stream's actual filename, as stored in the Streams2 Table 214, with the filename of such default message.

[0052] In an alternate embodiment, a script or other item of software running on the system generates an email that informs the client of the impending expiration date. For

example, an email message warning can be sent to the client a certain period (e.g., 7 or 30 days) before the expiration date, so as to avoid an unanticipated expiration of the client's content. In such an embodiment, the system repeatedly searches the database (e.g., CM Assets Table 212) to identify streams having an expiration date within the set period, places a corresponding entry in a queue or otherwise flags the content and, for each such stream, generates the email message.

[0053] Upon deletion of a stream, a record is also created in the Deleted Assets Table 244, specifying various stream information, as well as the expiration date and deletion date of the content.

[0054] In the event the stream to be deleted is currently being provided to an end-user 104, the system will be unable to delete the content. In such an instance, the system creates a record in the Deferred Delete Table 242, specifying the server name, file name and share information for the content to be deleted. As illustrated, for each unique combination of server name, share information and file name, the Deferred Delete Table 242 also stores the site ID, stream ID, the date on which the delete request was submitted to the system, the date of the last attempted deletion, and the number of failed deletion attempts. A script running on the system periodically searches the Deferred Delete Table 242 for entry and, upon the finding a populated record, proceeds to attempt to delete the content specified by the entry. Upon successful deletion, an entry is created in a Deleted Assets Table 244 and a status is changed in the Asset Locations Table 218.

Registration Process

[0055] Having described the components of the present embodiment, the operation thereof will now be described. As an initial matter, each client 102 must register with the service provider. In general, such registration includes the client 102 providing the service provider with account identifying information, a subset of which includes client identifying

information. More specifically, the client provides the account identifying and client identifying information contained within the CM Accounts Table 210. The client may also select the designed service level. In one embodiment of the present invention, the client 102 provides such information via a secure web page generated by the client web server 110. The client web server 110 receives the information and by executing a common gateway interface (CGI) script, writes the information via the LAN to the CM Database 118. In an alternate embodiment, the client 102 manually provides the information to the service provider, who, in turn, manually enters the information into the CM Database 118 via the administration terminals 114. In either embodiment, once the service provider receives the account identifying and client identifying information, a CM account ID is assigned and the corresponding record in the CM Accounts Table 210 is populated.

Client-Side Content Management Web Site

[0056] As noted above, the client web server 110 provides a content management web site that serves as an interface between the client 104 and the content management system. While the present embodiment utilizes such interactive web site for allowing the client 102 to interface with the system and to manage the client's content, it is to be understood that other interface devices may be used, including voice recognition devices, text-based systems, or by passing variables on a query string and running a script, or any other interface means. Furthermore, although the present embodiment passes data between the web server 110 and the script processor 116 in XML format, it is to be understood that other formats may be utilized. The content management web site of the present embodiment will now be discussed with reference to Figures 3a - j.

[0057] Once a client 102 logs into the system by submitting its username and password, the client is presented with the option to perform content management functions or administrative functions. As shown in Figure 3a, when the client 102 selects administrative

functions from a high level menu 302, the web page presents the client 102 with the ability to set account preferences, including default title, default author, default security key, and default security interval. As with the other web pages discussed herein, entry of information via the client 102 is received by the client web server 110 and placed in the appropriate fields in the CM database 118. With the account preferences web page of Figure 3a, the default preferences are stored in the CM Accounts Table 210. In alternate embodiments, the web page also displays account settings such as maximum allowable storage (max content Kb), actual storage used (e.g., sum of Kb size fields in Assets Table 212), and any other information stored in the database 118.

[0058] As shown in Figure 3b, when the client 102 selects the manage content option in the high level menu, the web site presents the client with six options: "Upload New Content", "Browse FTP Content", "Batch Uploads", "Playlists", "Search" and "Recycle". When the "Upload New Content" option is selected, the system presents the client 102 with the web page shown in Figure 3b. As discussed in greater detail below, the client 102 is able to identify a content file residing on its local machine, and upload the file to a repository server 106. In general, the client 102 specifies the file name and file details in an upload content form 306. These file details are then stored in the CM Assets Table 212 and the Streams2 Table 214. Notably, the account preferences set by the client 102 on the web page of Figure 3a are provided as default entries in the upload content form 306.

[0059] When the client 102 selects the "browse FTP content" option, the system presents the client with the web page shown in Figure 3c. As the name indicates, this option provides the client 102 with a listing of all content files previously uploaded by the client 102 to the client's FTP account on the FTP ingest server 108. The system first identifies the FTP ingest server 108 containing the client's FTP account. A call is made to the particular FTP ingest server 108, which in turn executes a script to read all file directories and files within

the client's account. The FTP server 108 returns a listing of such file directories and files in XML format to the web server 110. The web server 110, in turn, converts the XML information into HTML, which is displayed to the client 102. As shown in Figure 3c, the client 102 is presented with the name of its FTP account 308, as well as a listing of the client's file directories 310 and the client's content files 312.

[0060] The browse FTP content web page allows the client 102 to either delete a content file or cause a content file to be ingested from the FTP server 108 to a repository server 106. Such operations will be discussed in greater detail below.

[0061] The "batch uploads" option is shown in Figure 3d. As shown therein, the batch upload web page of the current embodiment is divided into two sections: one illustrating the batch files which have been uploaded to the client's FTP account, but not processed; and one identifying the client's batch files for which processing has begun 316. By way of example, batch file "1-100.bul" has been uploaded but not processed, while batch file "BillDefault.txt" was previously uploaded and processed.

[0062] It should also be noted that the batch upload web page may display any file detail stored in the CM database 118 corresponding to each batch file. More specifically, for each batch file listed as having been processed, information may be extracted from the submit time, process start time and process end time fields of the CM Batch Jobs Table 228.

[0063] The listing of available batch files 314 is generated by the system browsing the client's FTP content and extracting all pure text files. The listing of processed batch files is created by the system by searching the CM Batch Jobs Table 228 for all records identified by the client's CM account ID. The client 102 may select one of the available batch files for either deletion or processing.

[0064] The web page displayed when the client 102 selects the "playlists" option is illustrated in Figure 3e. In general, this web page provides the client 102 with the ability to

edit playlists. To this end, the web page presents the client 102 with a list of current playlists 318. The system generates the list of playlists by searching the CM Playlists Table 222 for all entries corresponding to the client's CM account ID, thereby resulting in a list of playlist or content group (CG) IDs, each of which identifies a different playlist for the client 102. For each CG ID, the system retrieves the CG description and format in the Content Group Table 224. The information is transferred in XML format and displayed in the playlist list 318 in HTML format.

[0065] The system also provides the user with a search form 320, whereby the client 102 can enter playlist details, such as playlist/CG ID, CG description and CG format, in which the system uses in searching the Content Group Table 224 in order to return a list of matching playlists.

[0066] The client 102 may also select a current playlist to manage. As shown in Figure 3f, when the client 102 chooses to manage a playlist, the system retrieves from the CM database 318 the CG id 322, various playlist details 324 as stored in the Content Group Table 224, and a list of streams or content files making up the identified playlist 326 as retrieved from the Content Group Streams Table 226 based on the CG ID. From the list of streams, the client 102 may select any number for deletion from the playlist, in which case the system removes the stream's id from the corresponding record in the Content Group Streams Table 226.

[0067] Notably, the web page also provides the client 102 with the uniform resource locator (URL) for the playlist 328. In general, the playlist URL 328 is incorporated into the client's web page as a link, and when activated by an end user 104, causes the playlist identified by the embedded playlist/CG ID to be played. The process of generating the playlist URL 328 and playing the streams associated with the playlist is described in greater detail with reference to Figures 5.

[0068] The web page displayed when the client 102 selects the "search" option is shown in Figure 3g. The system provides the client 102 with a search form 330 that enables the client 102 to enter the file details, including title, author, keywords, stream id, creation date, and duration, as search criteria. The web server 110 passes the file details entered by the client 102 to the script processor 116, which uses the values to search the CM Assets Table 212, CM Stream Keywords Table 216 and the Streams2 Table 214. The script processor 116, in return, returns various file details of the client's content files meeting the search criteria. Such results are shown in Figure 3h. The search form 330 also includes an option to display the results in XML format (as received from the script processor 116) and to search not only the client's current account, but also all child accounts (as indicated by the parent id field in the CM Accounts Table 210).

[0069] As an alternative to using the search form 330, the client 102 may simply browse all content in its content management account by clicking the "Browse" button. Once the client 102 clicks the browse button, the web server calls a script on the script processor 116 that retrieves all files associated with the client's username and password (and, thus, account id).

[0070] Exemplary search results are illustrated in Figure 3h. Each content file may be identified by any of the file details stored in the CM Database 118. In the present embodiment, each file is identified by the stream ID, title, stream type, stream description, creation date, file size, and status, all of which are stored in the CM Assets Table 212 and Streams2 Table 214.

[0071] Furthermore, the web page displaying the search and browse results (Figure 3h) provides the client 102 the option to select one or more files and either "delete" the selected files or "create a playlist" by adding the selected files to a new or existing playlist. In an alternate embodiment, clients 102 are also provided with the option to view the streams

URL playlist URL 328. In the event the "create a playlist" button is activated, the "create a playlist" window shown in Figure 3i is displayed. The window displays a list of the client's existing playlists 332 and the files contained within each playlist 334 (as provided in the "playlist" web page), and the window provides the client 102 a form 336 to enter playlist details, including description and format, for a playlist to be created using the files selected on the web page of Figure 3h.

[0072] It should be understood that the foregoing client web pages represent an exemplary client interface and that other interfaces are within the scope of the present invention. For example, clients may access their account via a script, such as a virtual basic script, on an active server page (ASP), and pass the relevant variable in a query string rather than entering them on a web page.

[0073] Thus, in certain embodiments clients 102 are able to manage playlists via an ASP page running on one of the system servers (e.g., the web server 210 or a separate application server, not shown) by passing certain variables on a query string. In one such embodiment, the client 102 log into the system by paring its username and password and preferably as part of the same query string, passes any of a number of parameters depending upon the desired function.

[0074] For example, the client may pass a parameter specifying what action is to be performed on the playlist. Such action may include: adding or creating a new playlist; editing the playlist (e.g., editing the playlist parameters or streams comprising the playlist); removing or deleting the playlist; listing playlist in the account (e.g., by name or ID); listing the playlists and the contents of each playlist and the like. The playlist ID parameter is specified to identify the playlist on which the system is to act.

[0075] By way of example, playlist parameter that can be specified (and thus edited) in the CM Database 118 is the name of the playlist. The playlist name can also be used to

identify the playlist which the system is to act. Another parameter that can be specified, (and thus edited) is the playlist format (e.g., RAM for RealMedia content and ASX for Windows Media content).

[0076] The client can also pass parameters specifying the action to be taken, if any, on individual streams comprising the specified playlist. Such actions may include, for example: adding a stream to the playlist; removing a stream from the playlist; removing all streams from the playlist; one or more stream IDs, to identifying the stream(s) when adding or removing them.

[0077] Preferably, under a client uses such an ASP to manage playlists, the server returns a message, such as an XML message, indicating whether the requested action was successfully performed and, if so, any information generated, much as the playlist ID of a newly created playlist. In the event the requested action was not successfully performed, the server responds with a error message identifying the error.

Uploading Content

[0078] As an initial matter, the system must identify the server to which a client's content should be uploaded. The site id is used to identify to which server a particular client's content should be uploaded, copied or moved. Specifically, once a client 102 logs into the system by providing its username and password, the system can identify the client's record in the CM Accounts Table 210, which provides the client's site id. Based on the site id, the system searches the CM Archive Volume Table 220 and identifies all records having the client's site id. Each of these records identifies a different server associated with the client's site id. If the system needs to access the FTP server 108 for the client (either to read the contents of or transfer a file to or from the client's account), the system identifies those CM Archive Volume Table records having the site id and a volume type corresponding to "FTP". Similarly, the system can identify repository servers 104 and streaming servers 120

associated with the client's site id by identifying those CM Archive Volume Table records having a type corresponding to "Repository" and "Streaming/Archive", respectively.

[0079] As described with reference to Figures 4a-b, the system of the present embodiment allows the client 102 to upload content to the repository servers 106 by any of several different processes. The first processes (described with reference to Figures 4a and b) transfer content files from the client's FTP account on the FTP server 108 to a repository server 106. As such, the first step in these processes is the client 102 uploading content to the FTP server 108 as noted above. The last manner is from the client's machine to a repository server 106 via HTTP upload.

[0080] As an initial step in the HTTP upload process of Figure 4b, the client 102 logs onto a web page provided by the client web server 110 by entering the client's user name and password. Having logged onto the system, the client then selects the "Upload New Content" option presented on the CM web page and identifies the content file locally stored on the client's machine that the client 102 desires to upload. Step 480. As illustrated in Figure 3b, the client 102 has the option of identifying the file by file name or by clicking the "Browse" button, which opens a window listing the client's locally stored files. Additionally, the client 102 specifies at least required file details for inclusion in the CM Database 118. Step 484. The client web server 110 proceeds to transfer the file from the client's machine onto the repository server 106 via an HTTP upload. Step 488. The system assigns a stream id to the file and creates a record in each of the CM Assets Table 212, Streams2 Table 214 and CM Stream Keywords Table 216, and the system assigns a location id and creates a record in the CM Assets Location Table 218. Step 492. The client web server 110 proceeds to write the file details, as provided by the client 102, into the appropriate fields in the CM Database 118. Finally, the system returns the stream id to the client 102, which is displayed on the CM web

page with a confirmation message. Step 496. In the event any of the foregoing steps fails, the system notifies the client 102 via a message on the CM web page.

[0081] Clients 102 may also programmatically upload content from their FTP Account on the FTP server 108 to a repository server 106 by executing a script. It is to be understood that although the following embodiment utilizes a Virtual Basic ("VBS") script on an active server page (ASP) of the web server 110, the functionality described may be implemented using any programming language such as Java Script.

[0082] As an initial step in the programmatic upload and ingest process, the client 102 uploads one or more files to the FTP ingest server 108 utilizing FTP client software residing on the client's machine. Once the files have been uploaded, the client 102 may use the ASP and pass certain variables on a query string to cause a file to be ingested or to perform another content management function.

[0083] As part of the ingest request of the present embodiment, the client 102 must enter its user-name and password, as well as certain file details. Required file details include the name of the file in the client's FTP directory ("Filename"), the bit rate of the file ("Bitrate"), the title of the file ("Title"), and the author of the file ("Author").

[0084] The following optional file details may also be provided by the client 102 depending on the desired action: the copyright notice to be added to the content file ("Copyright"), a description of the file ("Description"), key words associated with the file to be used in searches ("Keywords"), an indication as to whether or not the file is secure ("IsSecure"), an alphanumeric security string for use in accessing the file if it is identified as being secure ("SecKey"), the security interval, the description of an existing playlist in the client's account ("PlaylistDesc"), an indication whether to add the file to the playlist identified by the aforementioned description or to replace the existing files associated with the aforementioned playlist with the file being ingested ("PlaylistAction"), Playlist ID (CGID)

and an indication whether to automatically place the ingested file on a suitable streaming media server 120 or to only place the ingested file onto a repository server 106, where it will stay until the administrative staff manually causes the file to be placed on a streaming media server 120 ("AutoArchive"). In short, any stored file detail may be specified. The process of adding and replacing files in a playlist is discussed in greater detail below.

[0085] It is to be understood that the aforementioned file details are merely exemplary and that other file details may be considered required or optional in alternate embodiments of the present invention. Furthermore, it is to be understood such file details may be manually entered by the client 102 when submitting the ingest request, or, in an embodiment of the system wherein the client 102 enters the file details upon uploading the content files to the FTP server 108, the system may automatically retrieve the file details from the CM Database 118 based on the client's selection of the file to be ingested.

[0086] In general, the query string or command line identifies the active server page script, the location of the script, and the variables being passed to the script, including the user name, password, required file details, and any optional variables. The following is one example of such a command line:

```
http://contentmgmt.broadcast.com/vbs_script_ingest.asp?UserName=Foo&Password=Bar&
FileName=testfile.rm&BitRate=5600&Title=Test File&Author=John Doe
```

[0087] The foregoing exemplary command line passes the real media file "testfile.rm," which has a bit rate of 5600 kilobytes, is entitled "test file" and is authored by "John Doe." Furthermore, the client 102 is identified by the user name "Foo" and password "Bar." Such variables are passed to the script "Vbs_script_ingest.asp," located at "content mgmt.broadcast.com."

[0088] The following exemplary command line logs in user "Foo" using the password "Bar," uploading the file "Testfile," giving the file a Title, Author, Copyright, Description, and Keywords.

```
http://contentmgmt.broadcast.com/vbs_script_ingest.asp?username=Foo&Password=Bar&Titl
e=Test File&Author=Yahoo! Broadcast&CopyRight=(c)2000 All Rights
Reserved&Description= Testing The File System&KeyWords=test1|test2|test3&
```

[0089] When the script is called, the system receives the file details and causes the identified file to be moved from the FTP ingest server 108 to a repository server 106. Based on whether the file was successfully uploaded, the system provides to the client 102 one of two XML strings. Such XML string includes a success code field (scode), a success code description (scode_Description) and, if the upload was successful, the stream ID.

[0090] Thus, if the file was not successfully uploaded, the following XML string is returned to the client 102, wherein an scode of "1" indicates a failure and the value of the scode_description indicates the reason for the failure (e.g., an invalid username).

```
<?xml version="1.0" ?>
<response scode="1" scode_description="invalid username" />
```

[0091] In the event that the file was successfully uploaded, the system returns the following XML string wherein the `code` and `code_description` indicate the successful upload

```
<?xml version="1.0" ?>
<response code="0" code_description="Success" stream_id="999999" />
```

[0092] The system then renames the file and updates the appropriate fields in the CM database 118. More specifically, with the Stream ID assigned to the file, a record is created in each of the CM Assets Table 212, Streams2 Table 214 and CM Stream Keywords Table 216. Additionally, a location id is assigned and a record is created in the CM Assets Locations Table 218.

[0093] As will now be described with reference to Figure 4a, and continuing reference to Figure 3d, the client 102 may also perform two types of batch uploads, thereby transferring files from the client's FTP account to a repository server 106. The two types of batch uploads will be referred to as "default" and "standard" uploading. In general, default batch uploading applies default file details contained in the batch file to an entire sub-directory of content files. On the other hand, the standard batch upload process applies file details contained within the batch file on a file-by-file basis.

[0094] As an initial step, the client uploads the media files and batch file to the FTP ingest server 108. Step 446. In the case of a default batch upload, the batch file must be placed in the root directory of the client's FTP account. In the case of the standard upload process, the batch file, along with the media files, are left in the root folder of the client's FTP account.

[0095] The format of the batch file also depends upon the type of batch upload. The default batch upload file takes on the following format wherein the file is first identified as

being for use in a "DEFAULT" batch upload and the relevant file details are separated by double colons as follows:

```
DEFAULT::DIRECTORY::BITRATE::TITLE::AUTHOR::COPYRIGHT::  
DESCRIPTION::KEYWORDS::ISSECURE::SECURITYKEY:: SECURITY INTERVAL
```

[0096] The default batch file also includes an indication of the FTP DIRECTORY where the content files to which the default file details are to be applied are located. In the event the client 102 does not wish to supply an optional file detail, the location for such file detail is simply left blank. Furthermore, by placing the word "FILENAME" in parenthesis in either the Title or Description field in the batch file, the system will automatically include the file's file name in such field. By way of example, the following default batch file will apply to all files in the client's FTP directory named "288_Filesdirectory," which contains filing have a bitrate of 2880 kilobytes, entitled "Upload," including the file name, authored by J. Doe, a copyright notice reading "Copyright©2001," no description and no keywords, an indication that the file is secure, having the security key "mysecuritykey," and having a security interval of 122 minutes.

```
DEFAULT::288_FILES::2880::Upload(FILENAME)::J.DOE::Copyright©2001:::Y::MY  
SECURITYKEY::122
```

[0097] The default batch file also includes a PLAYLISTDESCRIPTION field. If the content that is being ingested needs to be added to an existing playlist, the existing playlist that it needs to be added to will be specified in the PLAYLISTDESCRIPTION field. Also included in the default batch file is a PLAYLISTACTION field. This field works in connection with the PLAYLISTDESCRIPTION field. When a stream is being ingested and a playlist is specified, the PLAYLISTACTION field will specify what actions are to be performed to the playlist. The actions that can be performed include: 1) adding the stream being ingested to the existing playlist whose name or description was specified in the

PLAYLISTDESCRIPTION field; and 2) replacing the streams that are already in the playlist specified in the PLAYLISTDESCRIPTION field with the new stream being ingested.

[0098] Also included in the default batch file is a SECUREAUTHIDREQUIRED field. This field is utilized if the client wishes to add an extra level of security to the content being uploaded. If this field is populated, an extra level of protection is added to the stream that is being ingested. Once this field has been populated, only users that are authorized to obtain the content will be granted access to it. The default batch file also includes an EXPIRATIONDATE field. As was discussed earlier, content on the system can expire and be deleted from the system at a specified time. This field allows the client to specify, at the time of uploading content, a date and or time at which the client wishes to have their content expire.

[0099] The standard batch upload batch file includes the same file details as the default batch file, however the standard batch upload batch file does not include the "DEFAULT" identifier or the "DIRECTORY" identifier. Instead, each line in the standard batch file begins with the file name to which the file details set forth in that line is to be applied. As such, the following is the format of the standard batch file:

FILENAME::BITRATE::TITLE::AUTHOR::COPYRIGHT::DESCRIPTION::KEYWORDS::I
SSecure::SECURITYKEY::SECURITYINTERVAL

[00100] Once the media files and the batch file is uploaded to the system, the client logs into the system via the CM website and selects the "Batch Upload" option, as shown in Figure 3d. Step 450. As part of the process of selecting the Batch Upload option, the client 102 is presented with a list of both processes and unprocessed batch files. (See Figure 3d). The client 102 selects an unprocessed batch file for processing.

[0100] Upon selecting a batch file for processing, the system executes a script that causes an entry to be created in the CM Batch Jobs Table 228. Step 454.

[0101] At some later point, according to a predefined interval, a task scheduler running on the script processor 116 causes a batch script to execute. Step 458. As described below, the batch script proceeds to read and process the unprocessed batch jobs identified in the CM Batch Jobs Table 228. In an alternate embodiment, processing of the batch files is not automated and instead is manually initiated by the administrative staff via the file management server 112.

[0102] Once the batch process is started, the system proceeds to read the first record in the CM Batch Jobs Table 228. Step 462. Specifically, the system reads the record to determine whether the start time field equals a null value or whether some actual start time. Step 466. If the start time field has an actual start time, then the batch process job has already been completed and the system proceeds to read the next record in the CM Batch Jobs Table 228. Step 462.

[0103] In the event that the start time value of the current record equals the null value, the system executes the batch job. Accordingly, the system causes the first file associated with the batch file to be moved to a repository server 106 and assigns the file a stream id. Step 470.

[0104] The system proceeds to rename the file that has been moved and updates the CM database 118 to reflect the associated file details provided by the batch file. Step 474. More specifically, the system parses the batch file based on the batch file format discussed above, inserting the file details into the CM Assets Table 212, Streams2 Table 214 and the CM Stream Keywords Table 216. The system also creates a record in the CM Asset Locations Table 218. Once the file details have been entered into the appropriate records in the CM database 118, the system determines whether or not all content files associated with

the batch file have been moved. Step 476. In other words, the system determines whether the end of the batch job has been reached. If the end of the batch job has been reached, then the system updates the record in the CM Batch Jobs Table 228 by entering the process end time. Step 478. If the end of the batch job has not been reached, then the system proceeds to move the next file associated with the batch file, assign a stream id, and update the CM database 118 accordingly. Steps 470, 474. Once the system cycles through the CM Batch Jobs Table 228, the processing of the batch jobs ceases until the scheduler begins the batch process again.

Selecting The Streaming Media Server on Which to Place a Client's Content

[0105] In the present embodiment, the system places each client's content on the streaming media server 120 that has the most free disk or storage space. More specifically, the script processor 116 (or other server in alternate embodiments) reads the client's record in the CM Accounts Table 210, identifying the client's site id. Using the client's site id, the system searches the CM Archive Volume Table 220 and locates all records having a volume type indicating streaming/archive. From this group of records, the system identifies only those having a platform (e.g., Windows Media or Real Media) corresponding to the format of the content file being uploaded. The resulting records are all potential streaming media servers 120 on which the content could be uploaded.

[0106] For each of the servers identified by the remaining records (each record and volume id corresponds to one server), the system proceeds to determine the available storage capacity by subtracting the used capacity from the maximum capacity. Using the records identifying all potential streaming media servers 120, the system notes the maximum storage capacity of each such media server 120, as set forth in the "max kb" field of the CM Archive Volume Table 220. The system then calculates how much of that capacity is actually being used by summing the size of all files stored on each such media server 120. More

specifically, for each volume id (server), the system identifies all of the records in the CM Asset Locations Table 218 having the same volume id and identifies the stream id's in each of those records. The result is a list of all files, as identified by stream id's, residing on the server 120 corresponding to the given volume id. The system reads and sums the files sizes by accessing the CM Assets Table 212. For each volume id, the system subtracts the total storage used by the files residing on the server having that volume id from the maximum storage capacity of the server, thereby calculating the actual available storage.

[0107] The content file being uploaded is placed on the streaming media server 120 (i.e., volume id) having the most available storage capacity. The system uses the selected volume id to create a new record in the CM Assets Locations Table 218 for the archived content file.

Creation of Playlist

[0108] As will now be described with reference to Figure 5, clients 102 may also programmatically create playlists in the CM Playlists Database. The client 102 logs into his/her account via the web server 110 and selects the "Playlist" link. Steps 510 and 515. The client 102 is presented with a form into which the client 102 enters playlist details, such as description and file format, and selects and orders stream files for inclusion in the playlist. Step 520. Once the client 102 has finished entering the playlist details, the system creates a new playlist record in the CM Playlists Database and assigns a CG ID to the playlist. This CG ID is returned to the client 102 for inclusion in the client's web site 102. Step 525. The system also creates corresponding records in both the CG Database and the CG Streams Database based upon the playlist details. Steps 530 and 535. Next, the system retrieves the stream ID for the next stream file in the playlist and enters the stream ID in a record in the CG Streams Database. Step 550. The system then queries whether the end of the playlist was reached. Step 555. If yes, then the system has completed creating records for the

playlist. However, if more streams are found, the system loops back to Step 550 to create additional records.

[0109] In this way, the playlist details are stored on the CM Database 118 so as to be accessible by the playlist server 122 to dynamically generate playlist redirector files, as described below.

Dynamic Playlist Redirector Files Creation

[0110] Referring back to FIG. 1, the media servers 120 are connected to the world wide web, or some other global communications network, through the LAN. In this respect, streaming content is made available to end users 104 through the world wide web.

[0111] Upon completion of the scheduling and production phase of the web-cast event, a uniform resource locator (URL) or link is returned to the client 102 to be embedded in the client's web page. An end user 104 desiring to listen to or view the web-cast on their computer or other device can click on the URL. A playlist or CG ID (12345) is preferably embedded within the URL, as shown below:

``

In the illustrative URL shown above, the link points to the playlist server 122 that will execute the "makeplaylist.asp" program and dynamically generate a playlist redirector file. One skilled in the art will recognize that although the "makeplaylist" application has an Active Server Page (or ASP) extension, it is not necessary to use ASP technologies. Rather, any programming or scripting language or technology could be used to provide the desired functionality, for example, .dll.. It is preferred, however, that the program run on the server side so as to alleviate any processing bottlenecks on the end user side.

[0112] The "makeplaylist" program functions to cause the playlist server 122 to make a call to the CM database 118 to retrieve the data necessary to dynamically generate the playlist redirector file. The playlist server 122 first accesses the CG Streams Table 226 and locates the record pertaining to the CG ID. As shown in FIG. 2, the CG Streams Table 226 contains the individual stream IDs and their respective sort orders. The stream IDs and the sort order are returned to the "makeplaylist" program which then makes a call to the Streams2 Tables 214 in the CM database 118 that contain data associated with the individual streams. Namely, the Streams2 Table 214 includes a URL prefix and a stream filename. The individual streams identifiers are also used by the "makeplaylist" program to call to the Stream-Servers Table 230 in the CM database 118. The Stream-Servers Table 230 includes the location or (DNS) hostname of the particular media server 120 containing the stream file associated with a particular stream identifier. Using the URL prefix, the hostname, and the stream filename, the "makeplaylist" dynamically generates a URL for each stream file. An example of such a URL is listed below, wherein the URL prefix is "mms://", the hostname is "mediaserver.company.com" and the stream file name is "filename.asf":

<"mms://mediaserver.company.com/filename.asf">

[0113] Using the individual stream URLs, the "makeplaylist" program then dynamically generates a redirector file to be passed to the media player stored on the end user's computer 104. An example of a redirector file for use with Windows Media Technologies (WMT) is shown below:

```
<ASX>
  <ENTRY>
    <REF HREF="mms://mediaserver.company.com/stream1.asf">
    <REF HREF="mms://mediaserver.company.com/stream2.asf">
    <REF HREF="mms://mediaserver.company.com/stream3.asf">
  </ENTRY>
</ASX>
```

[0114] One skilled in the art will recognize, of course, that different media technologies utilize different formats of redirector files and, therefore, the term "redirector file" is not limited to the ASX-type file shown above. Lastly, the end user's media player pulls each identified stream file from the media server 120 identified in the redirector file in the order in which it appears in the file.

[0115] It is to be understood that other system architectures are within the scope of the present invention. For example, certain embodiments include sites each having its servers on a separate LAN coupled to the administrative components (e.g., playlist server, file management server, script processor, database) via a secure SMTP connection. Furthermore, certain embodiments do not, as in the embodiment of Figure 1, include repository servers that maintain backed-up copies of files. One such embodiment that does not include repository servers for storing backups will now be described with reference to the schematic of Figure 6, the workflow diagram of Figure 7 and the CM Database schematic of Figure 8.

[0116] As illustrated in Figure 6, the architecture of the present embodiment essentially mirrors that of the embodiment of Figure 1. Indeed, the operation of the components illustrated therein is based on, and in many respects identical to the components described above with reference to Figure 1. However, rather than including repository servers for maintaining a backup of content stored on the streaming media servers, the present embodiment includes storage servers 606 that maintain a copy of each item of content to the exclusion of maintaining a duplicate copy on a streaming media server 620. In short, other than an external backup, such as on a tape backup system or mirror or redundant disk

system, preferably the only copy of each item of content for a given site is on the storage servers 606. Accordingly, as described in greater detail below, when providing access to the content, the streaming media servers 620 performs a LAN-based read of the appropriate content on the storage servers 606.

[0117] It should be appreciated that coupling multiple media servers 620 to a single storage server 606 having the only copy (vis-à-vis those media servers 620) provides several benefits. For example, the content management service provide saves money because less storage is used. Additionally, because multiple media servers 620 may access and, thus, stream the same item of content, clients 102 will encounter fewer instances of end users 104 not being able to receive the content because a media server 620 is already providing the maximum number of streams. While so utilizing the storage servers 606 and media servers 620 is preferable, it is within the scope of the present invention to have some sites employ an architecture wherein a repository or backup server and the media server each has a copy of the content, and have other sites employ a architecture where media servers do not have a copy of the content.

[0118] Unlike the embodiment of Figure 1, the present embodiment also includes site-specific file management servers 607. As described in greater detail below, each such file management server 607 causes files to be moved from each site-specific FTP ingest server 608 to the appropriate storage server 606.

[0119] It is also within the scope of the invention to have one or more sites in the present embodiment on separate networks in communication with the administrative servers and processors via SMTP or other connection, which themselves may be on a separate network.

[0120] Having described the architecture of the present embodiment, its general operation will now be described with reference to the workflow diagram of Figure 7. It is to

be understood, however, that aside from the differences noted or otherwise to be understood based on the following description, the components of the present embodiment operate as like named components of the embodiment of Figure 1. As illustrated therein, clients 102 are able to upload content either via an HTTP upload to the web server 610 or via an FTP upload to the client's primary site, as identified in the client's account record in the CM database 618. The web server 610, in turn, performs a LAN-based file move to the appropriate storage server 606. In alternate embodiments, the read is via a direct, point to point connection. Under direction of the file management server 612, each site-specific file management server 607 causes the content previously uploaded to the FTP ingest server 608 to be moved via the LAN to the appropriate storage server 606.

[0121] With the content stored on the appropriate storage server 606, each streaming media server 620 receiving a request for content will read the file from the appropriate storage server 606 (in the present embodiment, via a LAN based file read) to obtain the content and ultimately provide it to the end user 104 requesting it.

[0122] Because the functionality of the present embodiment, in large part, parallels that of the embodiment of Figure 1, the CM database 618 illustrated in Figure 8, largely parallels the CM database 118 of Figure 2, with similarly named tables and fields providing similar functionality and storing similar information unless stated or understood otherwise. However, the present CM database 618 includes variations to accommodate the use of storage servers 606, as well as to provide additional functionality. Such variations and additional functionality will now be described with reference to the database schematic of Figure 8 and continuing reference to the schematic of Figure 6.

Multiple Media Servers Per Storage Server

[0123] As noted above, an item of content stored on a storage server 608 is accessed by each media server 620 that provides an end user 104 access to such content. Furthermore,

in the present embodiment, multiple media server 620 may access a given item of content on a storage server 608. Consequently, the CM database 618 tracks which one or more media servers 620 are coupled to, and may access, a given storage server 608. To accommodate this potential one to many relationship, the CM database 618 includes the archive volume-media server ("AVMS") Table 846 and the Media Servers Table 848. Each record in the AVMS Table 846 associates a volume ID with a media server, as identified by media server (MS) ID, that is permitted to access content stored at the volume ID. Each MS ID, in turn, corresponds to a record in the Media Servers Table 848, which provides media server identifying information, as indicated. The "status" column stores one of the following values: CE (Check Enabled); MD (Manually Disabled); and ME (Manually Enabled).

[0124] As with the prior embodiment, each volume ID is essentially defined in a record in the Archive Volume Table 820. The volume type field in the archive volume table 820 specifies to which type of server (FTP ingest server, file management server, storage server, or web server) the volume ID corresponds. The server category field identifies the server as being used for business services or public services or any other category useful for reporting and maintenance purposes. The platform field is used to identify the media platform corresponding to the volume type or to indicate that the platform is not applicable or to indicate all platforms (e.g., a character such as "B").

[0125] The "archive server" field (in Archive Volume Table 820), and the "server name" field (in the Media Servers Table) store the computer's WINS/NETBIOS name for LAN-based file access (for example, PUB01), so that shared resources (such as directories, etc.) on the server can be accessed via that name (for example, \\PUB01\\PRoot1).

Account Categories

[0126] The present embodiment has various other features reflected in the CM database 618. For example, as illustrated in the Accounts Table 810, each account is

associated with an account category, as identified by account category ID (acct cat ID). Each account category ID corresponds to an entry in the Account Categories Table 850, which sets forth the name of the category. In general, an account category refers to a type of account, such as Internet radio or television or corporate account. The content management service provider may use such information for reporting and, because accounts can be segregated into broad categories, for simplifying the searching of accounts.

Time Zones

[0127] The Accounts Table 810 also includes a time zone ID field (TZ ID), which identifies a client-selectable time zone. Based on the client's selection, and the corresponding time zone ID, all information presented to the client 102 is based on the selected time zone. More specifically, the date and time field stored within the data base 618 are stored in a Greenwich meantime, however prior to being displayed to each client 102, the system converts the time into that of the client-selected time zone. Similarly, dates and times entered by each client 102 is converted to Greenwich Mean Time prior to being stored in the database 618. In the present embodiment, each time zone ID corresponds to a record in a time zone table (not shown), identifying the time zone and providing the relative difference between the Greenwich Mean Time and that time zone. Thus, by way of example, the expiration date, deletion date, creation date and any other time related information is presented in the Client's selected time zone frame.

Default Expiration Interval

[0128] The Accounts Table 810 also includes a default expiration interval field (default exp interval). In the event a client 102 fails to specify an expiration date for a particular file being uploaded, the system automatically adds the value stored (e.g., minutes hours, days, weeks, and the like) in the default expiration interval field to the date and time that the file is uploaded, thereby arriving at the default expiration date for the file. Such

default expiration interval may be set by the client 102 during registration, or at any other point in the process.

Accessors

[0129] The current embodiment also includes tables that specify who may access and alter account information. Specifically, the Account Accessors Table 852 and the Accessors Table 854 specify, for each account ID, one or more individuals (accessors) that may access the account. As illustrated, there may be several records in the account accessors table 852 for each account ID, thereby specifying multiple accessors for each account.

Stream Categories

[0130] Unlike the embodiment of Figure 2, the Streams2 Table 814 includes a field specifying each stream's category (category ID). Stream categories are essentially any type or division of content deemed helpful to the content management service provider or the clients 102. Furthermore, categories may be a broad or as focused as desired. In the present embodiment, stream categories are hierarchical in nature, with only the most detailed category being assignable to a stream. For example, the system may provide a first level category of "sports," with a second level of multiple subcategories, such as "baseball," "basketball," "football," and the like. Each subcategory, in turn, includes more detailed third level of subcategories, for example, under basketball, could be "NBA," "WBA" and "College." Only these most detailed categories are assignable in the present embodiment.

[0131] Preferably, each client 102 specifies a default stream category to be applied to uploaded streams that is stored in the Accounts table 810 (in the default stream cat ID field). When new content is uploaded via the client web site or programmatically, the client 102 has the option of using the default category or supplying a new category. In either case, the system creates a new record in the Streams2 Table 814 and populates the category ID field accordingly.

[0132] Each category ID corresponds to an entry in the Categories Table 838, which specifies the name of the category, the parent category of the subject category, the type of category (i.e., whether the category is assignable or not) and the level of the category in the category hierarchy. This system retrieves this information and constructs a visual display of the hierarchy on the client-side web page and permits each client to select the desired category. The Categories Table 838 is also keyed to the default category ID in the Accounts table 810.

Replication Sites

[0133] The present embodiment also allows each client 102 to have multiple sites – a primary site and one or more replication sites – associated with its account. In general, when uploading content via the client website, the client 102 may specify any one or more of the replication sites to which the content will be uploaded (in addition to the primary site). The client's primary site is identified by the site ID stored in the client's record in the Accounts table 810. The replication sites are identified in the Account Sites Table 856. More specifically, the Account Sites Table 856 includes records that associate multiple site IDs to the client's account ID. Of course, there is no requirement that a particular account have any replication sites associated therewith.

[0134] In operation, when a client 102 uploads content, the content gets passed to the storage server 606 associated with the client's primary site. As with the previous embodiment the stream ID assigned to the content gets returned to the client 102. Thus, in the context of the process of Figure 4b, an additional step 486 would be performed, namely, that if the client 102 is authorized to replicate its content to any site other than their primary site, the client optionally selects one or more replication sites to which the file being uploaded should be replicated.

[0135] In the background, when the client 102 uploads the content and identifies a replication site to which the content should be stored, an entry is made in a replication queue with the information identifying the content and the storage server 606 of the replication site. Periodically, a script checks the replication queue for entries and, for each entry in the queue, causes the appropriate content to be stored in the specified storage server 606 corresponding to the replication site. Thus, in the context of the process of Figure 4b, step 498 would be added to account for the system replicating the files to replication sites if one or more replication sites were chosen by the client while specifying file details. It should be noted that for each storage server 606 containing a copy of the content, there is a record in the Assets Location Table 818, each corresponding to the same stream ID and file name. This stream ID, universal to all copies of the file at the primary and replication sites, simplifies management of the copies, as there is only one entry in the Assets Table 812.

[0136] The ability to replicate content and the lack of a back-up copy of each item of content on a repository server, as in the embodiment of Figure 2, results in certain changes in the Asset Locations Table 818 status field. More specifically, the status is never "backing-up" and, instead, may indicate that the content is: available; being deleted; has expired; failed to be properly stored; being moved; new; being trashed; or being replicated, as described above. Notably, in the present embodiment, there is no need to copy a content file from one storage server 606 to another within the same site because multiple media servers 620 can access the content.

Recycle Bin

[0137] The status of being "trashed" is used in conjunction with a figurative "recycle bin," like that describe above with reference to Figure 3j. When browsing or searching content, a client 102 may select one or more items of content for deletion, or placement in the recycle bin. When selected for deletion, the system changes the value of the status field for

the corresponding record in the Asset Locations Table 818 to "trashed." The system also creates a record in the Deleted Assets Table 844 and enters as the deletion date a date (and time) some predetermined period in the future (e.g., one day, one week). A script running on the script processor or other server periodically searches (e.g., every three hours; every day; and the like) for records in the Asset Locations Table 818 having a status of "trashed." For each such record, the system locates the corresponding record in the Deleted Assets Table 844 and reads the deletion date. If the deletion date equals or is greater than (i.e., after the then current date (and time)), then the system proceeds to delete the content. Until such time as the content is actually deleted, the client 102 may restore the content from the recycle bin via a web page identifying the contents (or, in alternate embodiments, programmatically). When the client 102 selects an item of content for restoration, the system changes the status in the Asset Locations Table 818 from "trashed" back to "available" and deletes the record in the Deleted Assets Table 844.

Stream Groups

[0138] The present embodiment also allows clients 102 to logically group Streams into stream groups for ease of presentation to end users 104. More specifically, each logical stream group includes content files having the same susceptible content, but different media formats and/or bit rates. When a client 102 uploads content, the client web page presents the client 102 with the option of associating the uploaded content with an existing stream group or with a newly created stream group.

[0139] Stream groups are reflected in the Streams2 Groups Table 858, the Streams2 Group Table 860 and the Streams2 Group Streams Table 862. More specifically, the Streams2 Groups Table 858 contains a record for each Stream Group, as identified by Stream Group ID, associated with a particular account. The Streams2 Group Table 860 provides Stream information corresponding to each Stream Group, as identified by Stream Group ID.

Such Stream information includes, for example, the Stream Group name, title, author, description, associated key words, application which created, and/or type (e.g., multibit rate stream group or single bit rate stream group) of Stream Group. Lastly, the Streams2 Group Streams Table 862 identifies each Stream comprising a given Stream Group. Additionally, the Table 862 identifies the corresponding bit rate for the particular stream. By characterizing streams into groups, clients are able to easily track a particular piece of content and provide to their end-users 104 the stream having the appropriate media format and bit rate.

[0140] The client-side web site also preferably presents a web page for managing stream groups that parallels the page provided for managing playlists. For example, a stream group may be selected, identifying information, such as name, title, description, keywords and the like maybe edited, and the content of the stream group may be modified by adding or removing content (by adding or deleting, respectively, a record in the Stream2 Group Streams Table 862).

[0141] Alternatively, as with all content management functions, the stream groups may be edited via the client-side web site as described by passing variables in a query string to a script. In one such embodiment, the script requires values for three parameters: username-password; and stream group action, which may indicate: adding a stream to an existing group; creating a new group; editing a group; removing a stream from a group; listing the stream groups; or listing the groups and the contents thereof. Depending upon the action specified, the following optional parameters must be specified: name; title; description; author; keywords; miscellaneous, metadata of the group; whether promotable; stream group ID (used to specify a group to be edited or removed or to edit streams within a group); and stream action. Stream action, in turn, may indicate: adding a stream; removing a

stream; removing all streams; stream ID(s) (specified when adding or removing a particular stream).

[0142] The Stream2 Details Table 870 is related stream groups. Each record in the Stream2 Details Table 870 corresponds to a stream, as specified by stream ID. The information stored for the streams includes: (1) miscellaneous metadata, such as an extension of any other field that was truncated due to length of field restrictions; (2) if the stream supports multi bit rates, a stream group is created containing all bit rate versions of the stream; the stream group ID is stored in the multi bit rate stream group ID field (mbr sg id); and (3) an indication of whether or not the stream supports multiple bit rates (is multi bit rate).

[0143] Collectively, the information is used to identify those streams that support multiple bit rates and to identify the particular stream group (out of potentially many to which the stream belongs) corresponding to the stream's multi bit stream group.

[0144] At the time that the client uploads content, the default batch file also includes an optional STREAMGROUPNAME field. If the client wishes to include their content within an existing StreamGroup they can specify the Stream Group in this field. Thus any streams that are uploaded will then be added to the named Stream Group.

Stream Folders

[0145] The present embodiment also allows clients 104 to organize content into logical folders. The Streams Folders Table 864 and Folder Content Table 866 permit this. The Streams Folders Table 864 identifies folders, by folder ID, associated with each account ID. Each record also specifies the folder name and the ID of the parent folder, as folders may be hierarchical. In certain embodiments, the record also specifies metadata at a folder level that may be specified by the client or assigned by the system and applicable to the streams and/or stream groups in the folder. Such metadata includes, for example, name, type of

content, category, source of the content, and the like. The Folder Content Table 866 identifies the content within each folder, as identified by folder ID. More specifically, the present embodiment allows streams, stream groups, and playlists to be included within a folder. Consequently, the folder content Table 866 is designed to specify any of the foregoing types of content. To this end, the folder content Table 866 includes a content type field, which specifies to which of the foregoing types of content the record pertains. Each record also includes a content ID field which specifies the ID of the content to which the record pertains. Thus, if the content type specifies playlists, the content ID field will include the appropriate CG ID; if the content file specifies Stream Group, the content ID field specifies the Stream Group ID; if the content type specifies Stream, then the content ID field includes the Stream ID.

[0146] Preferably, the client web page includes a separate page that provides the client 102 the opportunity to both create folders and edit existing folders. To create a folder, the client 102 simply needs to provide the folder name and any parent folder to which the new folder relates and to select content comprising the folder. In essence, for each new folder created, a new record is created in a Streams Folders Table 864, and for each content added to the folder, a new record is created in a Folder Content Table 866. When editing an existing folder, the client web page preferably presents to the client 102 a listing of folders, as identified in the Streams Folders Table 864, and upon the client's selecting a folder, the system retrieves the details corresponding to the content ID stored in a Folder Content Table 866 and displays those details to the clients 102.

Promotability

[0147] The present embodiment also provides the clients 102 with the option of marking content as promotable. By marking content promotable, the client 102 essentially authorizes the content management service provider to select the content and place it on any

web site operated by the service provider. For example, if a client 102 provides news clippings, it may mark the content promotable, thereby authorizing the service provider to provide a link to the content on the service provider's news page. As illustrated in the CM database 618, Stream Groups may be identified as promotable in the Streams2 Groups Table 858 and folders may be identified as promotable in the Streams Folders Table 864. Clients 102 may indicate that a particular content is promotable when up-loading the content, or creating or editing a folder and stream group. In alternate embodiments, individual streams may be specified as promotable or not by including a promotable field in either of the Assets Table 812 or Streams2 Table 814.

[0148] The Database 618 also includes an FTP Info Table 868 for storing the site ID for each FTP Ingest Server 608, as identified by hostname. The Table 868 also identifies the username and password needed to access each FTP Ingest Server 608.

Playlist Resolution

[0149] Because the present embodiment does not store streams on the media servers 620 and because clients 102 can have replication sites, the processes of creating the playlist URL and creating the redirector file is different than described with regard to the embodiment of Figure 1. The differences will now be described with reference to Figures 6 and 8.

[0150] In the event a client 102 has one or more replication sites and the content to be presented is stored at multiple sites, the client preferably selects the content and the site from which the client desires it to be streamed. For example, a client may have a Japanese language stream; if stored on a site in Japan and a site in the United States, the client may desire it be streamed from the site in Japan. Consequently, the playlist URL of the present embodiment includes an indication of not only the content (e.g., stream ID, CG ID, which in the present embodiment are unique to each other), but also an indication of the location (e.g., site ID) from where the content should be streamed. An exemplary playlist URL, in which

the playlist server 622 is “playlist.company.com”, the program to dynamically generate the redirector file is “makeplaylist.dll”, the ID is “12345” and the site ID is “16”, is as follows:

`<href=“http://playlist.company.com/makeplaylist.dll?ID=12345&SITE_ID=16”>`

[0151] The playlist URL is constructed generally as described above. However, when the playlist is created, the system preferably provides the client 102 with a series of URLs, each corresponding to a different site at which the content is stored, and specifies the site name and time zone, as retrieved from the Sites Table 834. Alternatively, the client manually enters the site ID.

[0152] As with the previous embodiment, the makeplaylist component dynamically generates a redirector file based on the CG ID and, in the present embodiment, the site ID. More specifically, the component causes the playlist server 622 to access the Content Group Streams Table 826 and identify the streams within the identified playlist and their sort order. For each such stream, the server accesses the Streams2 Table 814 and retrieves the URL prefix and filename. As described in greater detail below, because the content is stored on the storage servers 606, and not the media servers 620, the filename refers each stream’s filename as stored on the storage server 606.

[0153] The playlist server 622 also accesses the Stream-Servers Table 830 to identify the hostname of the server on which each stream resides. It should be noted that because the client has a replication site with the stream, there will be multiple records in the Stream-Servers Table 830 for each stream ID. Consequently, the playlist server 622 must determine which record corresponds to the site identified in the playlist URL. The server does this by accessing the Servers2 Table 832 and identifying the record having both one of the

hostnames previously identified from the Stream-Servers Table 830 and the site ID from the URL. The hostname from this record is used in creating the redirector file. The system does this for each stream in the playlist.

[0154] In certain embodiment, a site will include multiple streaming media servers 622 capable of streaming a given content file. In other words, multiple media servers 622 can be coupled to a single storage server 606. In such an instance, the system must determine which of the media servers 620 to use. This many-to-one relationship will be reflected by multiple records in the Servers2 Table 832 having the appropriate site ID. In one embodiment, the decision is random or pseudorandom. However, in the present embodiment, the decision is based on load balancing and the weight field in the Servers2 Table 832.

[0155] The weight field is a preset indication of the percentage of streams to be served from a given media server 620, as compared to other media servers 620 capable of streaming the same content. For example, if two media server 620 were coupled to a given storage server 606 and could stream a file, and one could serve twice the streams of the other, then the weight of the higher capacity server would be approximately 66 and the weight of the lower capacity server would be approximately 33. When determining from which server to stream, the system would analyze the ratio of current streams being served of the two servers and utilize the server streaming a percentage of the total streams less than its weight. In other words, the system would seek to maintain the ratio of streams served equal to the ratio of the weights.

[0156] Once the media server hostname is determined, the component on the playlist server 620 obtains the filepath prefix from the record in the Stream-Servers Table 830 corresponding to the relevant stream id and hostname. As will be described in greater detail, the filepath prefix serves as a mounting point for the media server 620 so that the media

server 620 can access the storage server 606 to read the requested stream. This filepath prefix is appended to the beginning of the filename and placed in the redirector file.

[0157] An exemplary redirector file, in which the playlist comprises a single stream having a stream ID of “2000” and a filename of “/test/10/2000.asf”, the media server 620 hostname is “mediaserver.company.com”, and the mounting point is “proot1/PubShare01”, is as follows:

```
<ASX>
  <ENTRY>
    <REF
      HREF="mms://mediaserver.company.com/proot1/PubShare01/test/10/2000.asf">
    </ENTRY>
  </ASX>
```

[0158] It should be noted that in the present embodiment, the mount point is the same as the root share and share information for the server volume ID. Therefore, the media server 620 may read the correct contents of the storage server 606. Placing the filepath prefix in the redirector file is useful in situations where files stored on a storage server 606 are not allowed to be stored in one single share, but multiple shares are preferably created for the sake of efficiency of the backup and recovery processes (e.g., mirroring a portion of all shares). In this scenario, the hard drive(s) storage space on a storage server 606 is required to be stored under one of those shares.

[0159] By way of illustrative example, the Archive Volume Table 820 currently has the following record for a share created on the PUB01 storage server 606. The record in the Archive Volume Table 820 for that filer/share combination is as follows: cm_vol_id=142; volume_type=S; archive_server=\\PUB01; share=/PubShare01; server_addr=pub01.broadcast.com; platform=B; max_content_kb=131072000; status=CE; site_id=16; local_path=NULL; checking_status= [not used]; server_category=PUB; root_share=/proot1; attach_type=N; total_kb_used=90455873.

[0160] This would cause a record to be created in the Stream –Servers Table 830 having a filepath prefix of /proot1/PubShare01.

[0161] Let's say that the file path stored in the filename column in Streams2 table was '/test/10/2000.asf'. Assuming that the stream ID 2000 was a Windows Media stream that is accessed with the mms protocol, then playlist will construct the final URL as follows:

mms://mediaserver.company.com/proot1/pubshare01/test/10/2000.asf

[0162] In this scenario, for this URL to work correctly, the streaming server must have a mount point with the same name as the first component of the URL after the server hostname. In the above case example, the server name is mediaserver.company.com, and the first component of the URL after that hostname is "proot1." So, a mount point having the name proot1 is created on the streaming server such that it points to the following LAN file storage path: \\PUB01\proot1

[0163] A root share is optional, and is used only to reduce the number of mount points needed to be created on a streaming server in a situation whereby the storage device or hard drive on a storage server 606 has more than one share on it (for such reasons as efficiency of the backup processes).

[0164] A storage share on a storage server, however, is required because the systems will store a file in that share (further organized in account-specific directory, such as 'test' in the above example, and a hashed directory such as '10' in the above example). In the above example, the storage server 606 \\PUB01 has the following shares: proot1 and pubshare01.

[0165] Another embodiment of the present invention will now be described with reference to Figure 9. In the present embodiment, client 102 can upload content in a number of different ways, including to a central location or to a site-specific location. A central location is the location chosen by the system administrator that can be accessed from all other locations and is the location where various centralized servers are stored. For example the

system of the present embodiment includes Web Servers 910, which are located in a central location. The central location additionally includes a CM database 918 and Ingest Queue server 926, which issues commands to perform file movements, as will be discussed below in reference to Fig. 10. CM database 918 is similar to CM database 618 but includes additional tables and different fields, as will also be discussed in greater detail below. Content also can be uploaded to web servers 910 via HTTP ingest mechanisms.

[0166] Similar to Figure 6, the system also includes site-specific Storage Server 906, which stores the content and site-specific streaming servers 920 which stream the content. As noted above, the site may be a client's primary site which may be at a different location than the central location. Clients 102 can upload content to a site-specific storage server 906 via a site-specific FTP ingest server using file management server 907. The site-specific Streaming Servers 920 stream content to End User 104 via the Internet or other network.

[0167] It is to be understood that there can be numerous site-specific streaming servers 920 each of which can stream a different stream format such as the aforementioned formats. In one embodiment, a Server Classes Table is created in CM database 918 which stores a master list of all server classes based upon the stream formats. More specifically, each record in the Server Classes Table includes a unique server class ID, as well as the server class name and a description. This master list is used in connection with a Server Class Formats Table, which associates servers or a group of servers, as identified by their server class ID with a particular format. Thus, the Server Class Formats Table contains the server class ID as well as a stream format, thereby associating site-specific streaming servers 920 with a particular format (e.g., a server class having an ID "X" serves Real Media content). Such information may be used to identify the appropriate server when receiving, storing or streaming uploaded content.

[0168] Having described the architecture of the present embodiment, its general operation will now be described with reference to the workflow diagram of Figure 10. It is to be understood, however, that aside from the differences noted or otherwise to be understood based on the following description, the components of the present embodiment operate generally as like named components of the embodiment of Figure 6. As illustrated in Figure 10, clients 102 are able to upload content either via an HTTP upload to web server 910 or via an FTP upload to FTP Ingest Server 908 at the client's primary site, as identified in the client's account record in the CM database 918. After an FTP upload, the site-specific file management server reads the content from FTP Ingest Server 908, copies the content and then deletes the content from FTP Ingest Server 908. The site-specific file management server 907 then performs a LAN-based file move to the site-specific Storage Server 906, each of which are at the client's primary site as identified in CM Database 918. This is accomplished when Ingest Queue Server 926 sends a message or command to site-specific file management server 912 to move the file to site-specific Storage Server 906. Ingest Queue Server 926 then updates CM database 918 after the file move is complete. Similarly, if the content is uploaded to web server 910, web server 910 redirects the upload to the client's site specific file management server 912 located at the client's primary site which performs a LAN-based file move to FTP Ingest Server 908, updates CM database 918 and places an ingest message into Ingest Queue Server 926 to move the content to Storage Server 906. With the content stored on the appropriate storage server 906, each streaming media server 920 receiving a request for content will read the file from the appropriate storage server 906 (in the present embodiment, via a LAN based file read) to obtain the content and ultimately provide it to the end user 104 requesting it.

[0169] Ingest Queue Server 926 contains a message queue through which messages can be relayed to various sites within the system. These messages are placed into the queue

by web server 910 and are passed along to the site specific servers. In one embodiment, CM database 918 contains a CM Message Operations Table, within which is a list of potential operations to be performed on messages existing in the Ingest Queue Server 926. The CM Message Operations Table includes an operation ID as well as a name for the operation. Exemplary operations include, for example, ingesting content, moving content, copying content, deleting content, migrating content or replicating content. For example, an Operation ID of “001” may refer to ingesting content while an Operation ID of “011” may refer to replicating content.

[0170] In certain embodiments, the ability exists to have multiple operations queues which can handle specific operations. For example, one queue might handle deleting content operations while another queue might handle stream content operations. To facilitate this, the CM database 918 includes a table that stores details for the message queues. As such, each record in this table contains a message queue ID, queue name, a queue application field indicating the function of the queue, as well as a field indicating whether the queue is active. It should be appreciated that one advantage of this embodiment is that one operation with a higher priority will not be delayed in the queue while another operation with a lower priority is being performed. For example, if one client places a low priority request (e.g., to delete content) into the system, another client’s high priority request (e.g., to stream content) will not be delayed in the queue.

[0171] A particular account’s operations and/or a particular site’s operations can also be routed into a particular message queue. This has the benefit of always associating a particular account or site with a particular message queue, which may have certain specific parameters associated therewith. Thus, in one such embodiment having an account and a site queue, if an operation is requested (e.g., delete content), the system checks if the particular account associated with that content has a particular associated message queue. If the

account does have an associated message queue, the message is routed to that queue. If the account does not have a particular message queue, then the system checks if the site associated with that content has a particular associated message queue. If the site does have an associated message queue, the message is routed to that queue. If the account does not have a particular message queue, then the message is routed to a default queue. To effectuate this, the database 918 includes a table having both the operation ID, the Message queue ID, as well as the particular account ID. This table is populated at the time of creation of the account and associates the particular operation and message queue with the particular account. For example, when content associated with a certain account needs to be deleted, the system uses the message queue identified in the record corresponding to the account and the deletion operation. Similarly, the database 918 includes a table having both the operation ID, the Message queue ID, as well as the particular site ID. This table is populated at the time of creation of the site and associates the particular operation and message queue with the particular site.

[0172] It should be understood that an added result of using site specific streaming media servers such as 620 or 920 is that the streaming media servers can be located anywhere in the world and not necessarily geographically close to the central location (e.g., location of the system provider). For example, a client's primary site may be located in Japan while the central location might be located in the United States. This primary site may thus be located geographically close to the client's end-user, thereby reducing the time of delivery of content to the end-users.

[0173] In certain embodiments, to expedite the streaming of content, the system determines where a client's previous content has been stored. In one such embodiment, various additional tables in CM database 918 are utilized to accomplish this. For example, a Global Locations Table which is populated at the time that content is uploaded is used to

store identifiers such as the names of all of the worldwide site locations. As such, each record in the Global Locations Table contains a location ID, a location name, and the status of the site (e.g., whether or not actively receiving and/or streaming content). If the location is a subset of or otherwise related to another site, the table also includes a parent location ID so that it can be associated with its parent site. For example, there may be sites located in Taiwan, China and Japan that are each associated with a parent site location called Asia. It should be appreciated that a parent site location need not actually have an actual site and may just be a location in which related sites are located (e.g., Asia). When a stream is requested by an end user, the system first looks to the site that was specified by the end user when making the request. If the stream is not located at that site, the system will go to related sites at the same level. For example, if the content is not found in the Taiwan site, the system will proceed to look for the content at the China and Japan sites. If the system is unable to locate the content at any of the sibling sites the system will proceed to try and locate the content at a parent site (e.g., Asia). If the system is unable to locate the content at the parent site it will proceed to try and locate the content at sites that are at the same level as the parent (e.g., Asia's siblings). If the system is still unable to locate the content, the system will proceed to the grandparent site (e.g., Asia's parent). This process will proceed until all of the related sites have been searched. If the content is not found then the system will proceed to search all of the non-related sites. If the content is not found in any site, an error message will be returned to the user. The Global Locations Table is populated at the time that a particular site is utilized.

[0174] In order to properly manage all of the various global locations, a location ID field is added to the Servers2 Table when adding stream servers to the site. A location ID field is also added to the CM Sites Table in CM database 918 to associate a site to a particular location.

[0175] Additionally, the Accounts Table may also include a field that indicates that the account has been associated with a particular site at a particular location. As such, a location ID field can be added to the Accounts Table so that content that is uploaded relating to that account will be stored on servers at that location. This becomes the primary site for that particular account. This speeds up the uploading process by allowing the content to be uploaded to and eventually served from a server that is geographically closer to the client uploading or requesting the content.

[0176] In another embodiment, character sets that are supported by the system or by particular sites, servers or categories are identified in an additional table. This has the added benefit of providing an account with the ability to display the data in any language or form. This additional table associates a character set ID, a character set name and description, as well as a coding field to help decode the particular characters. The character set ID can then be included within records in the Accounts table to indicate the account's preference. For example, if a particular site or account mainly uses a foreign language with unique characters (e.g. Chinese or Japanese), and wishes to display the metadata (e.g., Title, Author) in that language, the character set ID field in the Accounts table would be populated with the value corresponding to that particular language. At the time the metadata needs to be displayed, the system would query the character set table for the particular coding for that language and any metadata that is served to that account will be displayed in that particular language.

[0177] In certain embodiments, clients are allowed to submit their content to the system either directly or encoded from a Batch Encoder system. According to one such embodiment, the client uploads content into Ingest Server 908 while grouping or categorizing this content into related content groups using XML files containing metadata information describing the content and how to relate or categorize the content. This metadata can include, for example, Title, Author, Date of Creation, Topic and any other information deemed

relevant to describing or identifying the Content. For example, a client may encode content by a particular news topic (e.g., Operation Enduring Freedom) and once the content is encoded it can be streamed by topic from Streaming Server 920 with greater ease.

[0178] To facilitate this XML uploading, the user uploads to Ingest Server 908, the content, as well as an XML batch ingest instruction file, which includes the information about the batch file including the metadata discussed above, as well as information about the format of the batch file. The file is then assigned a stream ID. If the batch file is not in a standard specified format, a custom XML transformation template written in Extensible Style Language Transformation (XSLT) can be created. To create this template, the system can create a map which allows the system to properly interpret the format of the batch file. For example, the system can create client XML tags which it will associate with XML tags that are already in the system. Once these tags are mapped to each other the system will be able to properly translate the XML batch file that the client has uploaded. An XSLT template field in CM Accounts Table stores the XML transformation template text which, the system accesses to interpret the metadata and/or content.

[0179] Within the CMS database of the present embodiment, numerous tables include fields to help identify and make use of the XML file received from the client. In the CM Batch Jobs Table, in addition to the other information about the batch job noted above, the stream ID that is assigned to the batch job XML feed is also stored in a stream ID field. This field is populated at the time that the file is assigned its stream ID. This allows the system to be able to identify the physical file for the batch job. Additionally, records in the CM Assets Table also include a CM Batch Job ID field, thereby associating the batch job with the item of content identified by a particular record.

[0180] In another embodiment, a CM batch job ID field is also included in the CM Streams2 Groups Table to store the ID of the XML Ingest Batch job during whose processing

a stream group was created, thereby associating stream groups with batch jobs. Similarly, a CM batch job ID field is also added to the CM Stream Folders table to capture the ID of the XML Ingest Batch job during whose processing a stream folder was created, thereby associating folders with batch jobs. It should be understood that once the streams, the stream groups and the folders have been properly associated with the batch jobs, metadata can be shared among them. This allows better interaction and characterization of all of the content.

[0181] In another embodiment, Streaming Servers can be grouped into categories. For example, the system is capable of assigning a group of streaming servers to a particular account. This is particularly useful, for example, if one account needs a relatively large capacity, (e.g., bandwidth) in which case the system assigns to the account one or more servers which are dedicated to a particular account. Conversely, if an account requires relatively few servers or little capacity, the system can allocate relatively fewer servers or a portion of a server's capacity to that account as well. Once this category has been assigned, the group of servers selected will be the only streaming servers used with this account. It should be understood that this categorizing need not be based on account and alternatively can be based on various other criteria such as location, stream format, or status of the client. An exemplary process through which servers are categorized will now be explained in greater detail.

[0182] At the time of account creation, a unique owner ID is generated and assigned to the account. This owner ID is stored, for example, in the CM Accounts table. A default streaming server category associated with the account's primary site as identified in the CM Sites Table is then associated to the owner ID that was just created. Additionally, an owner ID field may be added to the Stream Servers Table to specify the final streaming server category owner that the stream is delivered to. An Owner Category Info Table is also created at the time an account is created. More specifically, this table associates an owner ID, a

location ID, and a stream server category ID. This information is then associated with the particular account through the owner ID and it can be used to supply the owner's category to various other tables. For example, at the time of uploading content, an appropriate location ID is associated with the Stream ID in the Stream Servers Table. When a stream is requested, the system uses the owner id field and the location id field to locate the appropriate Stream Server Category ID. Once the appropriate Stream Server category is located, the content can be streamed from the correct server.

[0183] In one embodiment the Owner Category Info Table also stores information about the application that supplied the content. As was discussed above, client 102 can be an application that uploads content onto the system. A master list of all of the applications that upload content onto the system is stored in an Applications Table in CM database 918. This table includes an application ID as well as a field to store the name of the application. Such applications, may include, for example, web applications that upload content. This application ID can then be supplied to the Owner Category Info Table so that the system is aware of which applications are supplying which content, thereby allowing the system apply different rules to different applications. For example, content uploaded from a particular application may automatically be assigned to a particular streaming server.

[0184] An exemplary process of streaming categories will now be explained in more detail. If an account's content needs special streaming handling, for example because of the account needs a relatively large bandwidth, a new streaming server category is configured to contain a certain set of streaming servers capable of serving this account's content. These streaming servers are then assigned a streaming server category ID. This streaming server category ID is then associated with the owner ID corresponding to the account associated with the content requiring the special handling. After that, all streams uploaded for the account will automatically stream from the streaming servers in the streaming server category

associated with the owner that is attached to the account in question. To effectuate this, when content is uploaded, Ingest Queue Server 926 issues a command to store the content at a storage server 906 at the same site as the specified streaming server 920.

[0185] In order to facilitate the management of all of the streaming servers, CM database 918 includes a Streaming Servers Table which identifies all of the streaming servers available to the system. This Streaming Servers Table contains a Streaming Server ID as well as a server host name. This Streaming Server ID is then used by other tables to reference and associate data with the streaming server. For example, Server Category Streaming Servers Table utilizes the Streaming Server ID to keep a record of which servers are dedicated to which category by matching each Streaming Server ID with a Streaming Server Category ID. This facilitates the proper use of the correct servers.

[0186] To further facilitate the use of Streaming Server categories, the CM database 918 includes a Stream Server Categories Table which contains a master list of the streaming server categories. As such, each record in the table contains a Streaming Server Category ID, category name, description, as well as a location ID. Categories are normally hierarchical and can have subsets or different levels of categories within each other. Thus, the Streaming Server Categories Table may include information about any child, sibling, parent, grandparent or other related category to identify the relationship among different categories.

[0187] In another embodiment, content within the system that is requested frequently by users is copied or “cached” to the local storage drives of the streaming servers 920 from which it is served. This remedies the problem of where a minority of content on the system is requested much more frequently than the majority of the content on the system. It should be appreciated that the benefit of copying content to a local storage drive saves both time and resources by eliminating the transfer of files from the storage servers 906 to the streaming server 920 each time a request is made for that content. Content caching also allows the

system to not overburden the storage servers 906. An exemplary process of content caching in which the system automatically caches the content will now be described

[0188] In order to facilitate content caching, the system maintains a record of usage for each piece of content. More specifically, the system keeps track of how often a particular piece of content has been requested for streaming. This record is compared against a predetermined standard (for example, number of stream requests in a given time period or some other measure of use) to determine which content has been frequently requested. Alternatively, the records for each piece of content can be compared against each other to determine which has been requested more frequently. Once a piece of content is determined to have been requested often enough (e.g., more frequently than other content streamed from the server; streamed more frequently than the average frequency of other content on the system; more than a predetermined threshold; etc.) to require content caching, the system searches through CM Database 918 to determine which streaming servers are part of a streaming server category assigned to the particular account having that content that also matches the stream format of the piece of content. For example, if the piece of content belonging to a particular account has been designated to be “high demand” and needs to be streamed in Real Media format, the system locates all streaming servers that are associated with that account that stream Real Media content as determined based on information in the CM Accounts Table, Stream Server Categories Table as well as the Server Class Formats Table. After these servers are located, Ingest Queue Server 926 sends a command to File Management Server 912 to copy the content to the local hard drives of all of the appropriate streaming servers 920. CM database 918 is then updated by Ingest Queue Server 920 to reflect where the content has been copied to. In another embodiment there can be an additional server used for executing a script to initiate the movement of the content. It should be appreciated by one of skill in the art that content caching can be based upon criteria other

than demand. For example, content that is requested by a particular user may be copied to a local hard drive of a streaming server that particularly streams to that user.

[0189] It should also be understood that content caching can take place at any time after upload, even prior to being moved to the storage server. For example, a particular user may have his content copied directly to the Streaming Server 920 to facilitate quicker streaming.

[0190] Those skilled in the art will recognize that the method and system of the present invention has many applications, may be implemented in many manners and, as such, is not to be limited by the foregoing exemplary embodiments and examples. In this regard, any number of the features of the different embodiments described herein may be combined into one single embodiment and alternate embodiments having fewer than all of the features are possible. Moreover, the scope of the present invention covers conventionally known and future developed variations and modifications to the system components described herein, as would be understood by those skilled in the art.